generating a score associated with the lesion using a trained machine learning network.

11. An apparatus for classifying a lesion in a medical image, comprising:

means for receiving an input medical image depicting a lesion;

means for localizing the lesion in the input medical image using a trained localization network to generate a localization map;

means for classifying the lesion based on the input medical image and the localization map using a trained classification network; and

means for outputting the classification of the lesion,

wherein the trained localization network and the trained classification network are jointly trained.

12. The apparatus of claim 11, further comprising:

means for jointly training the trained localization network and the trained classification network by separately training the localization network to determine weights of the localization network during a first training phase, and training the classification network based on the weights of the localization network during a second training phase.

13. The apparatus of claim 12, wherein the means for separately training the localization network comprises:

means for receiving a multi-site dataset associated with different clinical sites and a deployment dataset associated with a deployment clinical site;

means for training a deep learning model based on the multi-site dataset; and

means for optimizing the trained deep learning model based on the deployment dataset to provide the trained localization network.

14. The apparatus of claim 11, further comprising:

means for generating a clinical relevance map and a label indicating whether the lesion is clinically significant based on the localization map and one or more of: a size of the lesion, an average intensity within the lesion, a variance of intensities within the lesion, radiomic features, and lexicon based features.

15. The apparatus of claim 11, further comprising:

means for generating a patch from the input medical image using the localization map; and

means for generating a score associated with the lesion using a trained machine learning network.

16. A non-transitory computer readable medium storing computer program instructions, the computer program instructions when executed by a processor cause the processor to perform operations comprising:

receiving an input medical image depicting a lesion;

localizing the lesion in the input medical image using a trained localization network to generate a localization map;

classifying the lesion based on the input medical image and the localization map using a trained classification network; and

outputting the classification of the lesion,

wherein the trained localization network and the trained classification network are jointly trained.

17. The non-transitory computer readable medium of claim 16, wherein:

the input medical image is a multi-parametric magnetic resonance imaging (mpMRI) image comprising a plurality of images;

localizing the lesion in the input medical image using a trained localization network to generate a localization map comprises localizing the lesion in each of the plurality of images using the trained localization network to generate a localization map for each of the plurality of images; and

classifying the lesion based on the input medical image and the localization map using a trained classification network comprises:

combining the localization maps for the plurality of images, and

classifying the lesion based on the plurality of images and the combined localization maps using the trained classification network.

18. The non-transitory computer readable medium of claim 17, the operations further comprising:

preprocessing the plurality of images of the mpMRI image to address variances between the plurality of images.

19. The non-transitory computer readable medium of claim 18, wherein preprocessing the plurality of images of the mpMRI image to address variances between the plurality of images comprises:

removing geometric variability in the plurality of images of the mpMRI image.

20. The non-transitory computer readable medium of claim 18, wherein preprocessing the plurality of images of the mpMRI image to address variances between the plurality of images comprises:

normalizing intensity variability in the plurality of images of the mpMRI image.

* * * * *